

NASA Contractor Report 187564

ICASE Report No. 91-40

ICASE

IMPLICIT SOLVERS FOR UNSTRUCTURED MESHES

V. Venkatakrishnan
Dimitri J. Mavriplis

Contract No. NAS1-18605
May 1991

Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, Virginia 23665-5225

Operated by the Universities Space Research Association

NASA

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23665-5225

(NASA-CR-187564) IMPLICIT SOLVERS FOR
UNSTRUCTURED MESHES Final Report (ICASE)

24 p

CSCL 01A

N91-24125

Unclas

G3/02 0019667

IMPLICIT SOLVERS FOR UNSTRUCTURED MESHES

V. Venkatakrishnan

Computer Sciences Corporation
Moffett Field, CA 94035

Dimitri J. Mavriplis

Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA, 23665

ABSTRACT

We develop and test implicit methods for unstructured mesh computations. The approximate system which arises from the Newton-linearization of the nonlinear evolution operator is solved by using the preconditioned GMRES (Generalized Minimum Residual) technique. We investigate three different preconditioners, namely, the incomplete LU factorization (ILU), block diagonal factorization and the symmetric successive over-relaxation (SSOR). The preconditioners have been optimized to have good vectorization properties. We also study SSOR and ILU themselves as iterative schemes. The various methods are compared over a wide range of problems. Ordering of the unknowns, which affects the convergence of these sparse matrix iterative methods, is also investigated. Results are presented for inviscid and turbulent viscous calculations on single and multi-element airfoil configurations using globally and adaptively generated meshes.

This research was partially supported under the National Aeronautics and Space Administration under NASA Contract No. NAS1-18605 while the second author was in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23665.

SECRET

CONFIDENTIAL

SECRET

CONFIDENTIAL

CONFIDENTIAL

CONFIDENTIAL

SECRET

1. INTRODUCTION

Impressive progress has been made in the area of algorithms for unstructured meshes in the last few years. Much attention has been focussed on improving the spatial discretization operator ([1-3]) which has evolved to a very high degree of sophistication. Usually explicit methods, such as Runge-Kutta schemes, have been used to march the solution to steady state. Some acceleration techniques such as local time stepping and residual averaging have also been implemented in this context. However, for large problems as well as stiff turbulent flow problems, the convergence rates of such methods degrade rapidly, resulting in inefficient solution techniques. In order to speed up convergence and propagate information more rapidly throughout the domain, more sophisticated multigrid or implicit methods are required.

The unstructured multigrid algorithm of Mavriplis [4] has been shown to produce efficient steady-state solutions for both the Euler and Navier-Stokes equations. In this approach, convergence acceleration is achieved by time-stepping on coarser unstructured meshes which may be generated independently from the fine mesh on which the equations are originally discretized. The principle behind this algorithm is that the errors associated with the high frequencies are annihilated by a carefully chosen smoother (a multi-stage Runge-Kutta scheme) while the errors associated with the low frequencies are annihilated on the coarser grids where these frequencies manifest themselves as high frequencies. The disadvantage of such an approach lies in the fact that the acceleration is achieved through the use of additional geometric constructions (i.e. user generated coarse meshes) which is often viewed as less desirable than for example an algebraic multigrid approach. A fully implicit method, wherein the system of linear equations is solved by direct methods, was developed and tested by Venkatakrisnan and Barth [5]. While providing a robust solution technique, direct methods are plagued by nonoptimal computational complexity and high storage requirements. Furthermore, for nonlinear systems with inexact linearizations, since the linear system of equations which arises at each time step need not be solved to a high degree of precision in order to maintain favorable overall (nonlinear) convergence rates, iterative implicit solvers may be employed.

Iterative implicit methods for unstructured problems have been investigated by Whitaker et. al [6], Hassan et al. [7], Struijs et al. [8] and Batina [9]. Venkatakrisnan [10] has tested preconditioned iterative methods on structured grid problems with special emphasis on vector performance issues. He concluded that some of these methods are quite competitive with other existing methods, while being readily applicable to unstructured grids. In this work we extend some of the ideas from [10] to unstructured grids.

Spatial discretization is achieved using piecewise-linear finite-elements. For dissipative terms, a blend of Laplacian and biharmonic terms is employed, the Laplacian term acting in the vicinity of shocks. The use of this particular discretization affords a relatively simple construction of the linear system, while enabling a straight-forward comparison of the implicit schemes with the previously developed multigrid strategy. For turbulent flow calculations, the unstructured mesh implementation of the Baldwin-Lomax algebraic model developed in [11] is incorporated. This model is not differentiable, and is therefore treated explicitly in the present scheme. The implicit methods investigated in this work are not restricted to any scheme in particular, and in the future may be applied to more complex upwind discretizations and more sophisticated multi-equation turbulence models.

2. IMPLICIT SCHEME

In non-dimensional conservative form, the full Navier-Stokes equations read

$$\frac{\partial w}{\partial t} + \frac{\partial f_c}{\partial x} + \frac{\partial g_c}{\partial y} = \frac{\sqrt{\gamma} M_\infty}{Re_\infty} \left[\frac{\partial f_v}{\partial x} + \frac{\partial g_v}{\partial y} \right] \quad (1)$$

where w represents the solution vector (conserved variables), and f_c and g_c represent the Cartesian components of the convective fluxes which are non linear functions of the w variables, and f_v and g_v are the Cartesian components of the viscous fluxes, which are functions of both the w variables, and the first derivatives of the w variables. The variables are stored at the vertices of a triangular mesh which is generated from a prescribed distribution of points by Delaunay triangulation [4]. Details of the spatial discretization using a finite volume scheme and its relation to a piecewise-linear finite element method may be found in [4].

The discretization of the governing equations in space leads to the following system of ordinary differential equations:

$$M \frac{dw}{dt} + R(w) = 0 \quad (2)$$

where R represents the spatial discretization operator, or the residual, which vanishes at steady-state and M represents the mass matrix, which contains the information relating the average value in a control volume to the values at the vertices. Since we are only interested here in steady state solutions, the mass matrix can be replaced by the identity matrix yielding

$$\frac{dw}{dt} + R(w) = 0 \quad (3)$$

If the time derivative is replaced by:

$$\frac{dw}{dt} = \frac{w^{n+1} - w^n}{\Delta t} \quad (4)$$

then an explicit scheme is obtained by evaluating $R(w)$ at time level n , and an implicit scheme by evaluating $R(w)$ at level $n+1$. In the latter case, linearizing R about time level n , one obtains:

$$\left(\frac{I}{\delta t} + \frac{\partial R}{\partial W} \right) \delta W_i = -R_i \quad (5)$$

$$\delta W_i = (W^{n+1} - W^n)_i$$

Eqn. (5) represents a large nonsymmetric linear system of equations for the updates of the vector of unknowns and needs to be solved at each time step. As δt tends to infinity, the method reduces to the standard Newton's method. The term $\frac{\partial R}{\partial W}$ symbolically represents the implicit side upon linearization and involves the Jacobian matrices of the flux vectors. The discretized convective fluxes are linearized exactly on the left-hand side of the equation. Only a first order accurate representation of the artificial dissipation terms is employed in the linearization on the left hand side, due to storage considerations. This results in the graph of the sparse matrix $\frac{\partial R}{\partial W}$ being identical to the graph of the supporting unstructured mesh (i.e. every vertex in the matrix is connected only to its nearest neighbors). The sparse matrix thus has a symmetric structure, even though the matrix itself is not symmetric. Linearization of the complete biharmonic dissipative terms would result in a much denser matrix with a different graph, since

each vertex would also be connected to its second to nearest neighbors. The storage requirements for the representation of such a matrix become prohibitive. The penalty in making this approximation in the linearization is that we can never approach Newton's method (with its associated quadratic convergence property) due to the mismatch of the right and left hand side operators in Eqn. (5). The viscous fluxes are linearized with a few approximations. First, the laminar viscosities, which are computed using Sutherland's law, are not linearized in the energy equation, and the average quantities at the cell centers are approximated as well. The validity of these approximations has been established by solving a very low Reynolds number laminar flow at very high CFL numbers (non-dimensionalized time steps). Second, the algebraic turbulence model, being nondifferentiable is not linearized and is treated explicitly.

Since the linear system is itself approximate there is little to be gained by solving it to a great precision. To obtain favorable overall (nonlinear) convergence, it has been found that it is better to solve the linear problem to a moderate degree of precision and proceed to the next time step. However, for stiff problems it may well be necessary to solve the linear problem well and one has the control to do so in the present framework. The time step in Eqn. (5) is taken to be inversely proportional to the L_2 norm of the residual. Since we have a mismatch of operators in Eqn. (5), it is necessary to limit the maximum time step.

The system of linear equations is solved in the present work by the GMRES technique developed by Saad and Schultz [12]. There is a host of iterative methods for solving nonsymmetric linear systems. Each of these methods has its own advantages but in the present context we shall just employ one: GMRES. Venkatakrisnan [10] compared the Chebychev semi-iteration technique to GMRES for structured CFD problems and found GMRES to be marginally better. Moreover, the choice of a particular iterative technique is not as important as that of a good preconditioner; and the better the preconditioner, the more computationally intensive it is, diminishing the relative importance of the iterative method. Without a good preconditioner, most of these iterative methods fail to converge for the kind of stiff problems which arise in computational fluid dynamics.

The GMRES technique is quite efficient for solving sparse nonsymmetric linear systems and is outlined below. Let x_0 be an approximate solution of the system

$$A x + B = 0 \tag{6}$$

where A is an invertible matrix. The solution is advanced from x_0 to x_k as

$$x_k = x_0 + y_k$$

GMRES(k) finds the best possible solution for y_k over the Krylov subspace $\langle v_1, Av_1, A^2v_1, \dots, A^{k-1}v_1 \rangle$ by solving the minimization problem

$$\|r_k\| = \text{Min}_y \|v_1 + A y\|$$

$$v_1 = A x_0 + B \quad , \quad r_k = A x_k + B$$

GMRES procedure forms an orthogonal basis v_1, v_2, \dots, v_k (termed search directions) spanning the Krylov subspace by a modified Gram-Schmidt method. Storage is required to store these search directions. As k increases, the storage increases linearly and the number of operations, quadratically. To mitigate this, Saad and Schultz also describe GMRES (k,m) which is a restarted GMRES (k), where the k search directions are discarded and recomputed every m cycles. GMRES can also be thought of as an optimal polynomial acceleration scheme. Preconditioning greatly improves the performance of GMRES as well as the other related iterative

methods. It decreases the size of the spectrum so that the optimal polynomial generated by GMRES can better annihilate the errors associated with each eigenvalue.

3. PRECONDITIONING

Instead of Eqn. (6) the preconditioned iterative methods solve the following systems:

$$P A x + P B = 0 \quad (7)$$

$$A Q (Q^{-1} x) + B = 0 \quad (8)$$

The systems of linear equations in Eqn. (7) and Eqn. (8) are referred to respectively as, left preconditioned and right preconditioned systems and P and Q as left and right preconditioners. The role of the preconditioner is to cluster the eigenvalues around unity. For reasons given in [10] we shall just employ right preconditioning. We have examined three preconditioners, namely the incomplete LU factorization, SSOR and block diagonal. We will describe below the preconditioners and the optimizations done to extract the best vector performances out of them.

A simple choice is a block diagonal preconditioner which computes the inverse of the 4x4 diagonal block associated with a grid point. Good vectorization when using this preconditioner is easy to achieve by unrolling the LU decomposition of the 4x4 diagonal matrix as well as the forward and back solves over all the grid points. A family of preconditioners arises out of an incomplete LU factorization and is referred to as ILU(n). Here n represents the level of fill-in. n=0 implies no fill-in beyond the original nonzero pattern. In the present work ILU(0) is used since it is quite robust and has lower storage requirements. It is also possible to cast the symmetric successive over-relaxation (SSOR) as a preconditioner as has been shown by Saad [14]. Saad recommends setting the relaxation factor to 1 when using SSOR as preconditioner. In this case the SSOR preconditioner looks exactly like the ILU preconditioner, except that the lower and the upper factors are read off directly from the matrix A rather than by an incomplete factorization. The incomplete factorization is a nonvectorizable procedure (although parallelizable by using wavefront ordering described below) and SSOR preconditioning dispenses with this sequential procedure. We will also test ILU factorization and SSOR as iterative techniques by themselves for solving the linear sub-problems at each time step.

4. DATA STRUCTURES

In this section we describe the data structures and kernels employed which are critical in reducing memory requirements and obtaining good performance. In the course of the GMRES method with preconditioning as per Eqn. (8) we need to address two kernels.

The first kernel is a sparse matrix - dense vector multiply to compute $A x$. The most commonly used data structures [15] are not ideal for this purpose since they have poor vectorization properties. The ITPACK data structure, which allocates storage based on the maximum number of nonzeros in a row, is inefficient for sparse matrices arising from unstructured grids, because the degree of a vertex is arbitrary. The data structure that we use for storing the sparse matrix A is most easily explained by interpreting the underlying triangular mesh as an undirected graph. Associated with each edge are the two vertices, say n1 and n2, which are incident to the edge. The spatial discretization operator (the right hand side) utilizes this data structure and therefore, this information is already available. We store the two 4x4 matrices which contain the influence of n2 on n1 (entry in row n1 and column n2 in A) and vice versa. The diagonal blocks are stored separately. With such a data structure, we can carry out a

matrix vector multiplication efficiently by employing a coloring algorithm to color the edges of the original mesh to get vector performance. Note that the data structure deals with blocks of 4x4 matrices; for a scalar matrix the above mentioned data structure is roughly equivalent to the coordinate storage scheme [15]. However, since the graph of the sparse matrix is equivalent to that of the supporting unstructured mesh, the matrix is known to have a symmetric structure (although the matrix itself is not symmetric). Hence, we achieve a savings with respect to the standard coordinate storage scheme by only storing the coordinates of the upper half of the matrix.

The second kernel deals with the effect of the preconditioner Q on a vector. Q is D^{-1} for block diagonal preconditioning and $(\tilde{L} \tilde{U})^{-1}$ for ILU/SSOR preconditioning, where the $\tilde{\cdot}$ indicates approximate factors. The block diagonal case is straight-forward in this aspect and was discussed earlier. The ILU/SSOR preconditioners require repeated solutions of sparse triangular systems. By using a level scheduling (also known as wavefront ordering) [16,17] it is possible to obtain good vector performance. Under this permutation of the matrix, unknowns within a wavefront are eliminated simultaneously. The key step in this procedure is an off-diagonal rectangular matrix - vector multiply. This requires that \tilde{L} and \tilde{U} be stored in a convenient form and we choose a data structure similar to that of A . In addition to the nonzero blocks and the column numbers which are provided by the factorization, we store the row numbers. With this additional information, the data structure becomes similar to the edge-based data structure employed for the A matrix except that we only store one block per edge. The off-diagonal matrix vector multiply can then be vectorized by interpreting the rectangular matrix as a directed graph and coloring the directed edges. The performances are further enhanced by performing all the operations on blocks of size 4x4 since we are dealing with coupled systems.

The memory requirements for the present algorithm are linear in n , the number of vertices. The implicit scheme requires three arrays of size $7 \times 16n$ in addition to a few integer arrays of size n . One of these arrays stores the matrix A in the edge-based data structure, a second in the YSMP format which is suitable for the factorization and the third contains the \tilde{L} and the \tilde{U} factors. The factor 7 comes from having 3 times as many edges as vertices (valid for all 2-D triangular grids, neglecting boundary effects); we store two blocks per edge plus the diagonal matrix for all the vertices. The second array is reused for storing the search directions in GMRES, permitting up to 27 search directions to be stored. Block diagonal preconditioning dispenses with one of these arrays.

The ordering of unknowns has a bearing on the convergence properties of many iterative methods. This is true for iterative methods which involve a directional bias such as the SSOR/ILU preconditioning. For structured meshes in [10,18] it was found that a column-major ordering which minimized the bandwidth (the "most local" ordering) yielded the best convergence rates. For unstructured meshes we have settled on the Reverse Cuthill-McKee (RCM) ordering [15]. This is a standard ordering used in sparse direct methods to reduce fill-in, but it also appears to be the "most local" ordering. We have also tested orderings based on coordinates of the vertices (sorting the vertices by the x coordinates, y coordinates or some combination of x and y coordinates). The RCM ordering gives marginally better convergence rates over a wide range of problems. RCM is also more efficient in that it creates fewer wavefronts, thus producing longer vectors.

To achieve good overall vector performance, careful attention also needs to be paid to the assembly of the matrix. In the present set-up, the matrix assembly is performed by looping over the edges as far as possible. This is easily done for the inviscid fluxes and the first order

dissipative terms, but is quite involved for the full viscous fluxes. We have found it expedient to assemble the matrix for the viscous fluxes by looping over the triangles instead and coloring the triangles to achieve vectorization. The Jacobians are derived analytically, but with some approximations for the viscous terms as was discussed earlier.

5. RESULTS AND DISCUSSION

The iterative method outlined above requires a few parameters. The start-up CFL number and the maximum CFL number that can be used need to be specified. It is also possible to freeze the factorization after a few time steps (or after a prescribed reduction in the residual) and increase the efficiency of the code, since it eliminates the assembly and/or the factorization of the matrix. This is an additional parameter. GMRES requires a few parameters. It requires the maximum number of search directions k , the number of restart cycles m and a tolerance level which specifies the desired order of reduction of the residual of the linear sub-problem. The solution to the linear system is terminated when the number of iterations exceeds the specified maximum whether or not the tolerance criterion is met. In all the problems, the tolerance is set to 10^{-5} .

We first study a standard airfoil case, namely inviscid flow over the ubiquitous NACA0012 airfoil at a freestream Mach number of 0.8 at 1.25° angle of attack. The unstructured grid contains 4224 vertices or 8192 triangles. A close-up of the nearly uniform grid is shown in Fig. 1a. The solution (not shown here) agrees with standard results. We obtain lift, drag and moment coefficients of 0.3523, 0.0226 and -0.0452 respectively. The convergence histories of five different methods are shown in Fig. 1b as a function of CPU time. Since we are dealing with different methods which require varying amounts of work at each time step we believe that CPU time is the only true measure for comparing them. Since there are quite a few parameters involved in each of these methods, what we have shown is the "best" convergence history obtained with each method. GMRES with ILU preconditioning (GMRES/ILU) uses 5 search directions, CFL $20 \cdot 10^6$ and freezes the factorization after 30 time steps. GMRES/SSOR, wherein SSOR is used as the preconditioner, employs 15 search directions, CFL $20 \cdot 10^6$ and freezes the matrix after 30 time steps. GMRES/DIAG, which uses block diagonal preconditioner, employs 25 search directions with 3 restarts, CFL $10 \cdot 500,000$ and freezes the preconditioner after 25 time steps. The ILU iteration uses CFL 1-50 and freezes the matrix after 25 steps. Finally, the SSOR iteration uses CFL 1-25 and freezes the matrix after 30 time steps. Using multiple "inner" sub-iterations with the ILU and the SSOR iteration schemes in order to be able to use larger time steps turns out to be less efficient for this problem. The number of time steps taken by GMRES/ILU, GMRES/SSOR, GMRES/DIAG, ILU and SSOR are 75, 100, 75, 700 and 700 respectively. The parameters given above for the five methods, we believe, are nearly optimal for this problem and yield the best convergence history for each of the methods. Having to choose many parameters is a major drawback in using iterative methods to solve the approximate linear systems arising from nonlinear problems. However, we will be able to provide some guidelines for choosing these parameters for the best of these methods, namely GMRES/ILU, by solving a few more representative problems. In Fig. 1b, we notice that GMRES/DIAG is quite slow even for this simple problem, while ILU iteration appears to be quite good. SSOR iteration and GMRES/SSOR have similar convergence histories. SSOR as a preconditioner is not as effective as the ILU preconditioner; GMRES/ILU appears to be the best of all the methods. As we shall see, as the problems get bigger and more stiff, GMRES/ILU performs much better than the other four methods.

We next consider inviscid subcritical flow over a 4 element airfoil at a freestream Mach number of 0.2 and angle of attack of 5° . The triangular mesh employed has 10395 vertices. The grid is shown in Fig. 2a. The solution is not shown here and may be found in Mavriplis [4]. In Fig. 2b we present the convergence histories of GMRES/ILU, GMRES/DIAG and ILU and SSOR iteration. GMRES/SSOR had great difficulties in the initial stages and is not shown. GMRES/ILU converges much better than the other methods. The parameters for GMRES/ILU are 10 search directions and CFL $20 \cdot 10^6$, the factorization being frozen after 30 time steps. GMRES/DIAG employs 25 search directions with 2 restarts, CFL $10 \cdot 5 \cdot 10^5$ and freezes the preconditioner after 30 time steps. ILU iteration uses CFL 1-50, freezes the matrix after 50 time steps and does not use sub-iterations. SSOR iteration uses CFL 0.5-5 and freezes the matrix after 100 time steps. The number of time steps taken by GMRES/ILU, GMRES/DIAG, ILU and SSOR are 100, 70, 400 and 400 respectively. SSOR, either by itself or as a preconditioner, is clearly unsatisfactory.

We compare the performances of the methods on a transonic turbulent flow over an RAE2822 airfoil, referred to as Case 6. The flow conditions are $M_\infty = 0.729$, $\alpha = 2.31^\circ$ and Reynolds number 6.5×10^6 based on the chord. The flow is computed on a mesh with 13751 vertices which contains cells in the boundary layer and the wake region with aspect ratios up to 1000:1. The grid is shown in Fig. 3a. The pressure plot and skin friction distribution and experimental data are shown in Figs. 3b and 3c. The lift, drag and moment coefficients are 0.7342, 0.0132 and -0.0978. Fig. 3d shows the convergence histories of the various methods. We notice that only GMRES/ILU and GMRES/DIAG converge, the latter doing so much more slowly. GMRES/SSOR diverges for any reasonable CFL numbers at all and its convergence history is not shown. The parameters for GMRES/ILU are 25 search directions and CFL 5-25000. We freeze the factorization after 80 time steps. We also freeze the turbulence model after nearly six orders of reduction in the residual; otherwise, the residual hangs and the convergence of the method slows down. The effect of freezing the turbulence model in this fashion has minimal effect on the aerodynamic coefficients (less than 0.02% change in lift coefficient). The parameters for GMRES/DIAG are the same as for GMRES/ILU. The number of time steps taken by both GMRES/ILU and GMRES/DIAG is 150. The unstructured multigrid algorithm of Mavriplis [4] takes nearly 300 secs. on the YMP to reduce the L_2 norm of the residual to $.3 \times 10^{-3}$ and GMRES/ILU takes about 450 secs. to get to the same level (7 orders of reduction in residual) for this problem. In the full multigrid algorithm, the problem is first solved on coarser grids, whereas GMRES/ILU starts from freestream conditions on the fine grid. The ILU and SSOR iterations use 10 sub-iterations, CFL .5-2.5 and still do not converge after 200 time steps.

The final case computed is turbulent flow over a four-element airfoil computed on an adapted grid with 48691 vertices. The grid and a close-up view near the leading edge are shown in Figs. 4a and 4b. The flow conditions are $M_\infty = 0.1995$, $\alpha = 16.02^\circ$ and Reynolds number of 1.187×10^6 . The convergence histories with and without freezing the turbulence model are shown in Fig. 4c as a function of the CPU time. The number of time steps taken is 400. The multigrid algorithm takes 2100 secs. to reduce the residual to 1.79×10^{-2} while GMRES/ILU takes about 2000 secs. to reach the same stage (five orders of reduction of the residual). The computed Mach contours for this case are shown in Fig. 4d, illustrating the complexity of this flow. In Fig. 4e the computed surface pressure distribution is compared with experimental wind-tunnel data.

In summary, we have found that for inviscid flows 5-10 search directions are usually sufficient whereas the turbulent viscous cases require 25 search directions with GMRES/ILU. The start-up CFL number is usually about 20 for inviscid problems and about 5 for turbulent viscous cases and the CFL number is allowed to increase up to 500-50000 fold. We use non-restarted GMRES whenever possible, which eliminates one of the parameters and is better suited for stiff problems (see [12]). The GMRES/ILU runs at about 90-120 MFlops on the Cray YMP (uni-processor) at the NAS facility, with performance improving as the problems get larger.

6. CONCLUSIONS

We have compared five candidate implicit methods for solving the compressible Navier-Stokes equations. For inviscid problems, with a small number of vertices and low cell aspect ratios, many of the methods work well, GMRES with ILU preconditioning performing the best. For larger problems, especially at high Reynolds numbers, almost all the methods except for GMRES/ILU converge extremely slowly, if at all. Not surprisingly, SSOR, either as an iteration or as a preconditioner, suffers dramatically as the problem increases in size or in the degree of complexity. GMRES/ILU is quite competitive with the unstructured multigrid algorithm, while eliminating the need for independent coarse grids to be generated. It does, however, incur a larger memory overhead than the multigrid algorithm. Even though these methods have been compared for a particular spatial discretization, we believe the trends should hold for other discretizations as well. We have carried out a number of optimizations to extract the best vector performances out of all these methods. Finally, the turbulence model itself appears to inhibit convergence in the latter stages. This needs further investigation and perhaps incorporating a field equation model with proper linearization would solve the problem.

ACKNOWLEDGEMENTS

The authors would like to thank T.J. Barth of NASA Ames Research Center for his contribution in arriving at the edge-based data structure for sparse matrices. The first author also thanks the NAS Applied Research branch at NASA Ames Research Center for supporting this project.

REFERENCES

1. Jameson, A., Baker, T. J., and Weatherill, N. P., "Calculation of Inviscid Transonic Flow over a Complete Aircraft", *AIAA Paper 86-0103*, January, 1986.
2. Stoufflet, B., Periaux, J., Fezoui, F., and Dervieux, A., "Numerical Simulation of 3-D Hypersonic Euler Flows Around Space Vehicles Using Adapted Finite Elements", *AIAA Paper 87-0560*, January 1987.
3. Barth, T. J., and Jespersen, D. C., "The Design and Application of Upwind Schemes on Unstructured Meshes" *AIAA Paper 89-0366*, January, 1989.
4. Mavriplis, D. J., "Multigrid Solution of the Two-Dimensional Euler Equations on Unstructured Triangular Meshes", *AIAA Journal*, Vol 26, No. 7, July 1988, pp. 824-831
5. Venkatakrishnan, V. and Barth, T.J., "Application of Direct Solvers to Unstructured Meshes for the Euler and Navier-Stokes Equations Using Upwind Schemes", *AIAA Paper 89-0364*, January, 1989.

6. Whitaker, D.L., Slack, D.C. and Walters, R.W., "Solution Algorithms for the Two-dimensional Euler equations on Unstructured Meshes", *AIAA 90-0967*, Reno, NV, January, 1990.
7. Hassan, O. , Morgan, K. and Peraire, J., "An Implicit Finite-Element Method for High Speed Flows" *AIAA Paper 90-0402* January 1990.
8. Struijs, R., Vankeirsblick, P. and Deconinck, H., "An Adaptive Polygonal Finite Volume Method for the Compressible Flow Equations", *AIAA 89-1959-CP*, Buffalo, 1989.
9. Batina , J.T., "Implicit Flux-Split Euler Schemes for Unsteady Aerodynamic Analysis Involving Unstructured Dynamic Meshes", *AIAA Paper 90-0936* April, 1990.
10. Venkatakrishnan, V., "Preconditioned Conjugate Gradient Methods for the Compressible Navier-Stokes Equations", *AIAA 90-00586*, Reno, NV, January, 1990. To appear in *AIAA Journal*.
11. Mavriplis, D. J., "Algebraic Turbulence Modeling for Unstructured and Adaptive Meshes", *AIAA Paper 90-1653*, June, 1990.
12. Saad, Youcef and Schultz, M.H., "GMRES: A Generalized Minimum Residual Algorithm for Solving Nonsymmetric Linear Systems", *SIAM J. Sci. Stat. Comp.*, Vol. 7, No. 3, 1986, pp. 856-869.
14. Saad, Youcef "Basic Iterative Methods", *Supercomputing 90*, New York, November 1990.
15. Duff, I.S., Erisman, A.M. and Reid, J.K., "Direct Methods for Sparse Matrices", *Oxford Science Publications*, Clarendon Press, Oxford, 1986.
16. Anderson, E. and Saad, Y., "Solving Sparse Triangular Systems on Parallel Computers", *International Journal of High Speed Computing* Vol. 1, No. 1, May 1989, pp. 73-96.
17. Saltz, J.H., "Automated Problem Scheduling and Reduction of Synchronization Delay Effects", *SIAM J. Sci. Stat. Comp.*, Jan., 1990.
18. Duff, I.S. and Meurant, G.A., "The Effect of Ordering on Preconditioned Conjugate Gradients", *BIT*, Vol. 29, 1989, pp. 635-657

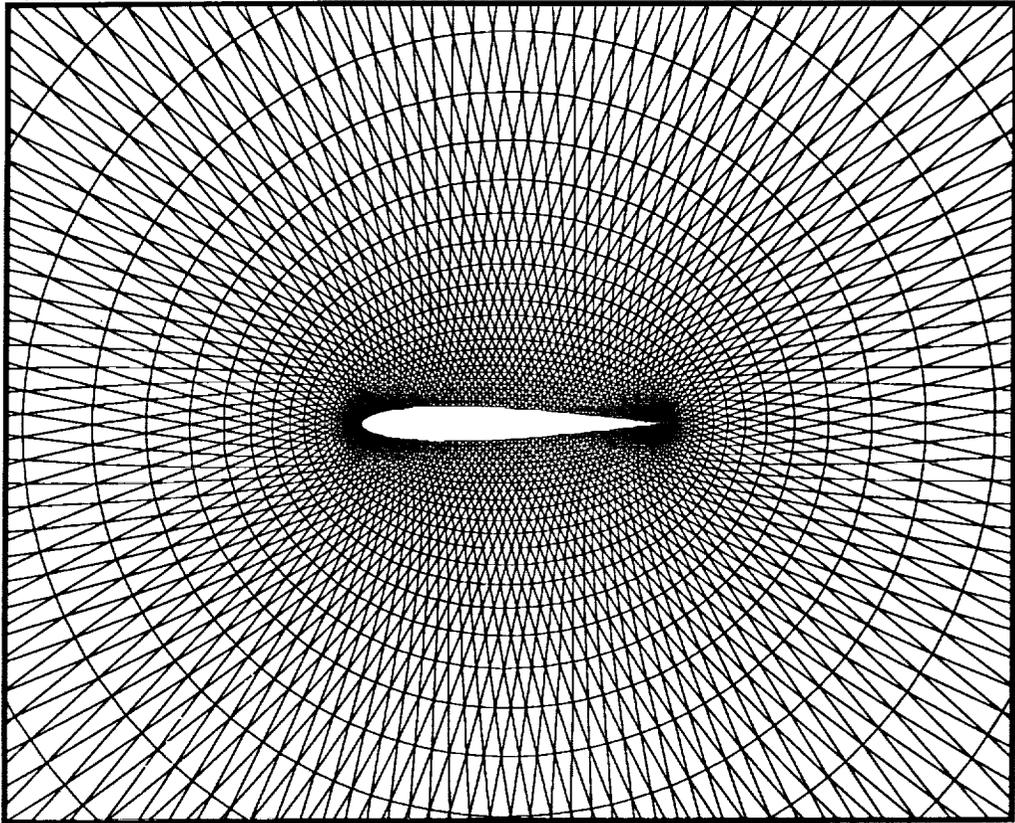


Figure 1a
Mesh for Computing Inviscid Flow over NACA 0012 Airfoil
(Number of Vertices = 4224)

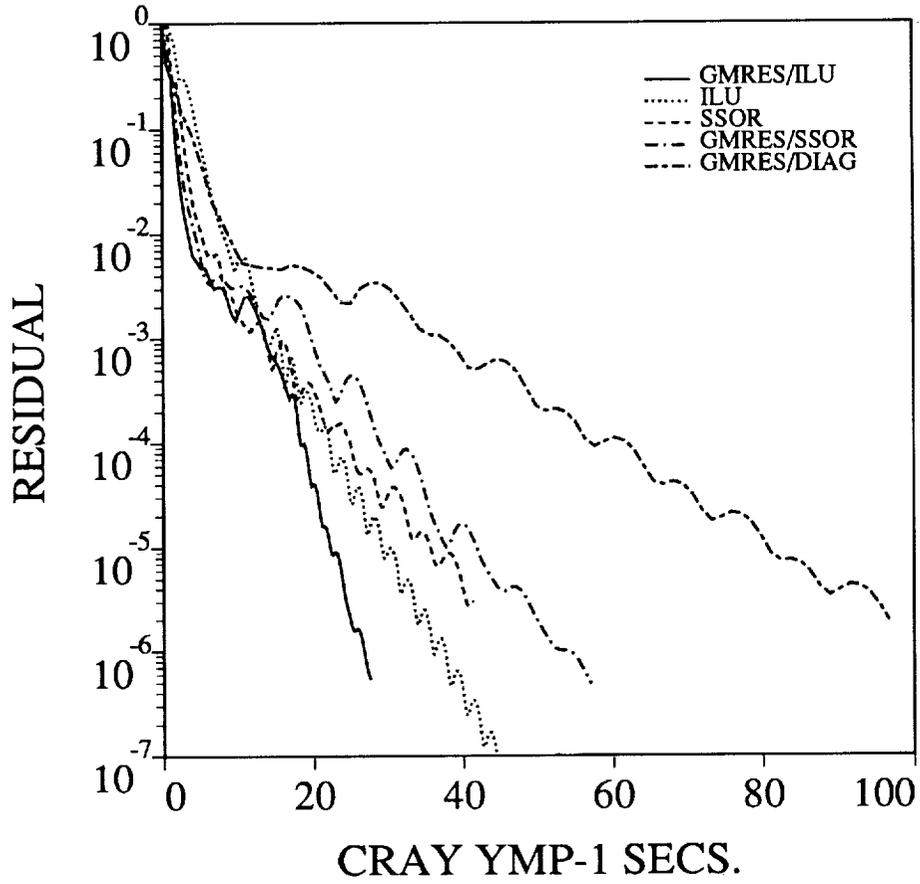


Figure 1b
Convergence Histories of Various Implicit Methods
for Inviscid Flow Over a NACA 0012 Airfoil
(Mach = 0.8, Incidence = 1.25 degrees)

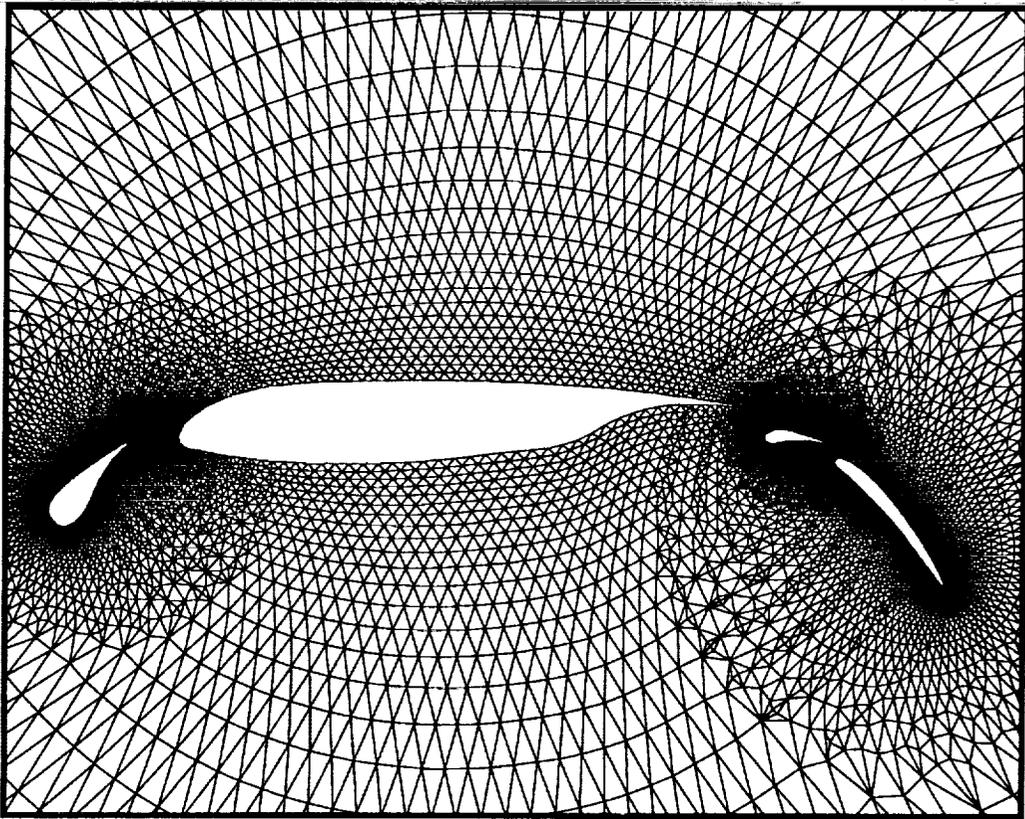


Figure 2a
Mesh for Computing Inviscid Flow over a Four-Element Airfoil
(Number of Vertices = 10,395)

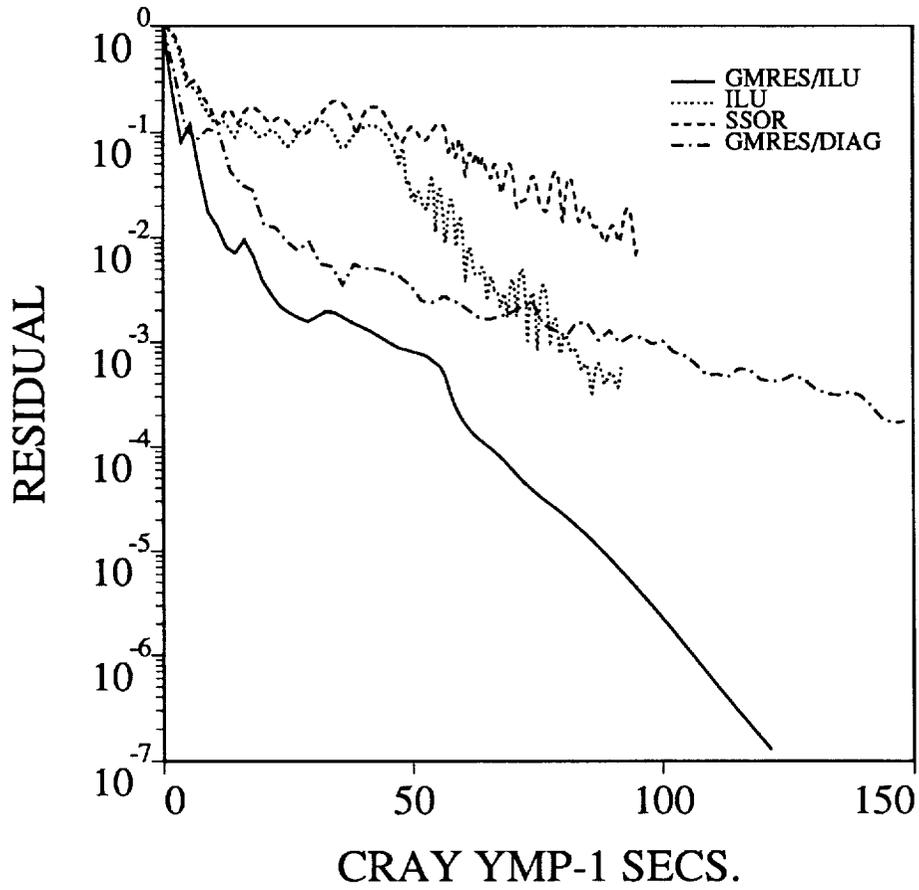


Figure 2b
Convergence Histories of Various Implicit Methods
for Inviscid Flow Over Four-Element Airfoil
(Mach = 0.2, Incidence = 5 degrees)

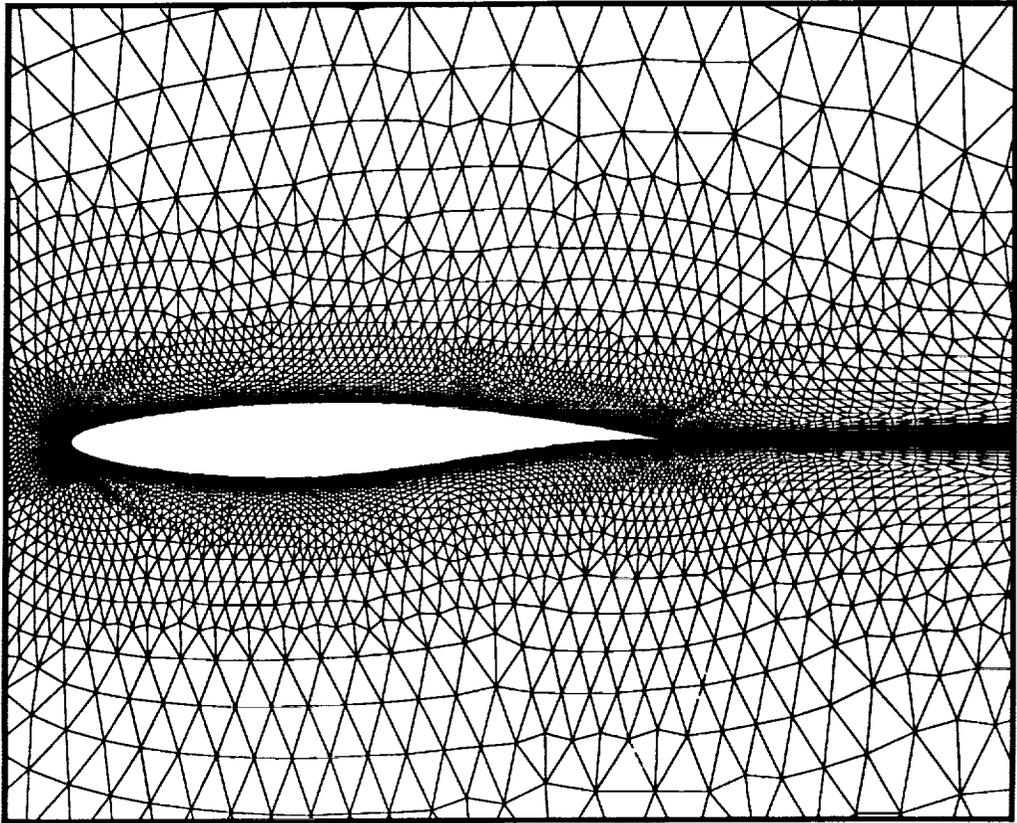


Figure 3a
Mesh for Computing Viscous Turbulent Flow Over an RAE 2822 Airfoil
(Number of Vertices = 13751)

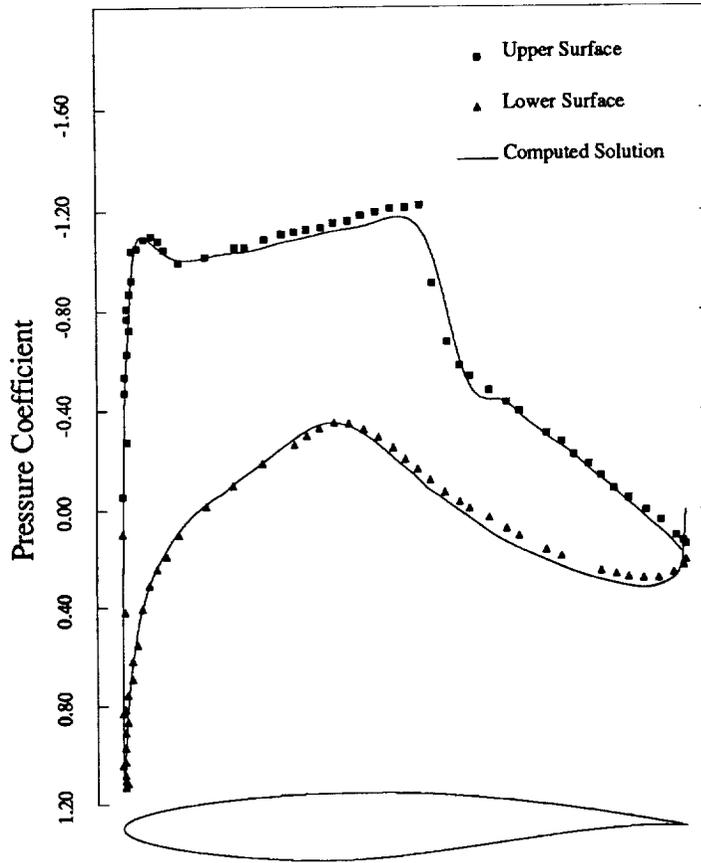


Figure 3b
Computed Surface Pressure Distribution for Viscous Turbulent Flow over RAE 2822 Airfoil
(Mach = 0.729, Incidence = 2.31 degrees, Reynolds Number = 6.5 million)

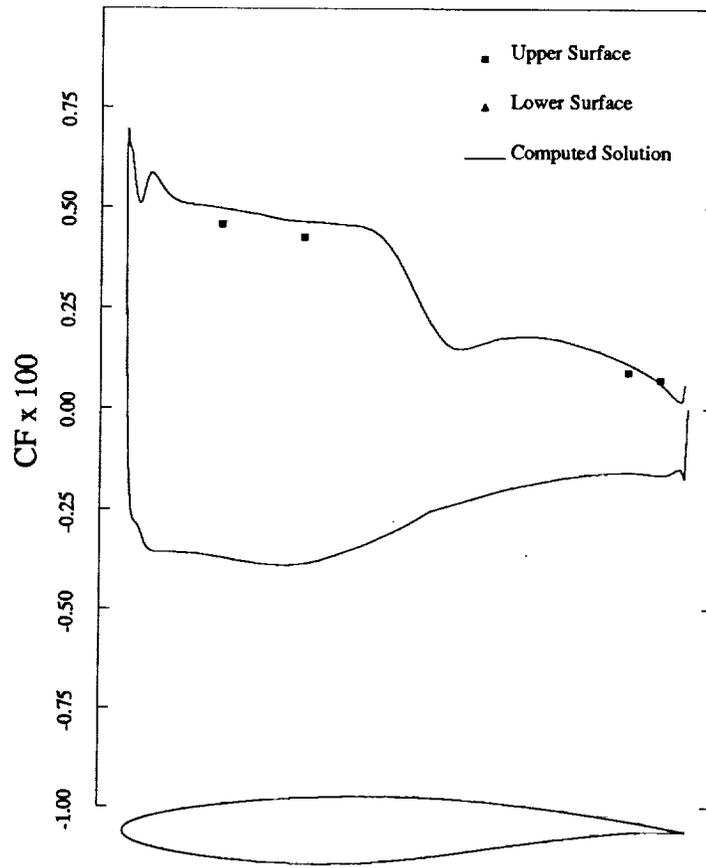


Figure 3c
Computed Skin Friction Distribution over RAE 2822 Airfoil
(Mach = 0.729, Incidence = 2.31 degrees, Reynolds Number = 6.5 million)

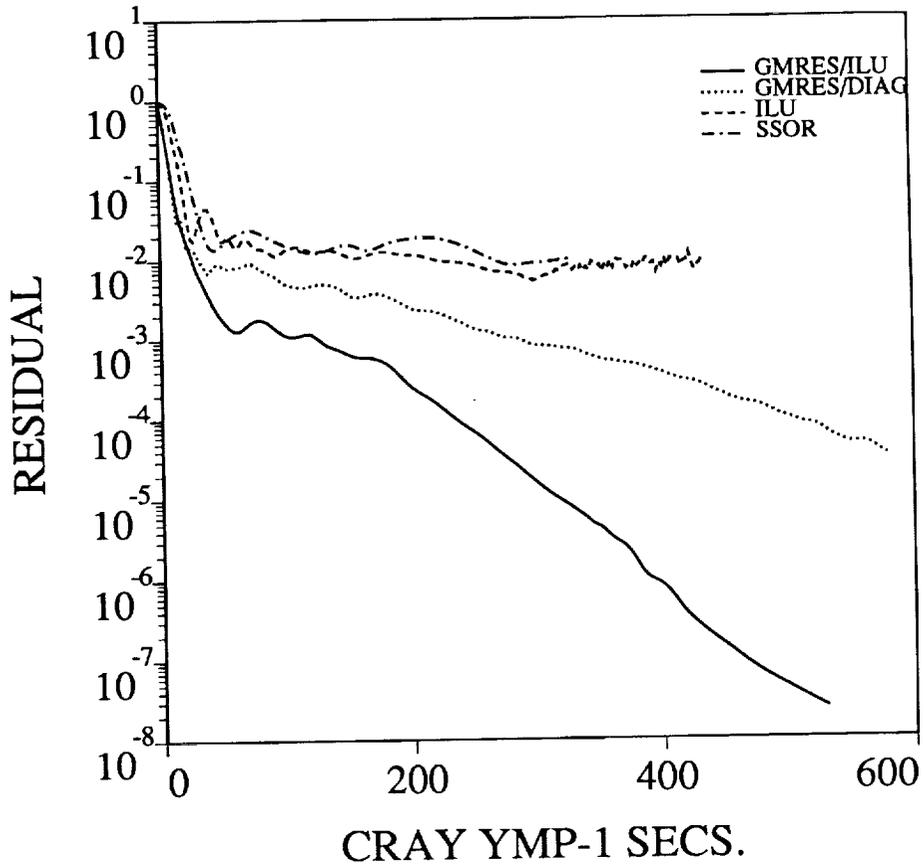


Figure 3d
Convergence Histories of Various Implicit Methods for Viscous Turbulent Flow
over an RAE 2822 Airfoil
(Mach = 0.729, Incidence = 2.31 degrees, Reynolds Number = 6.5 million)

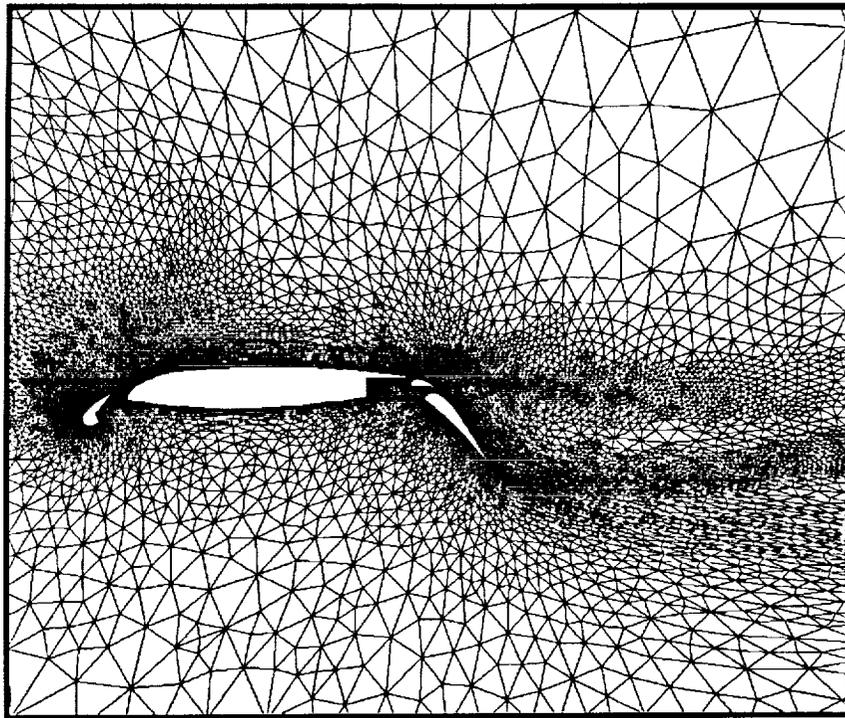


Figure 4a
Global View of Adaptively Generated Mesh for Computing Viscous Turbulent Flow over a Four-Element Airfoil
(Number of Vertices = 48,691)

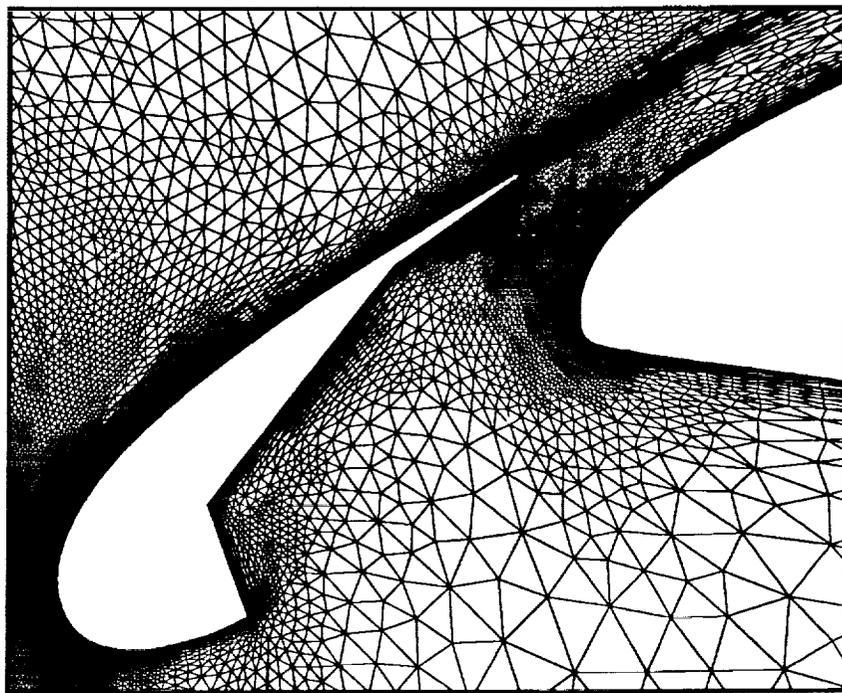


Figure 4b
Close-up View in Region of Leading Edge of Adaptively Generated Mesh about Four-Element Airfoil

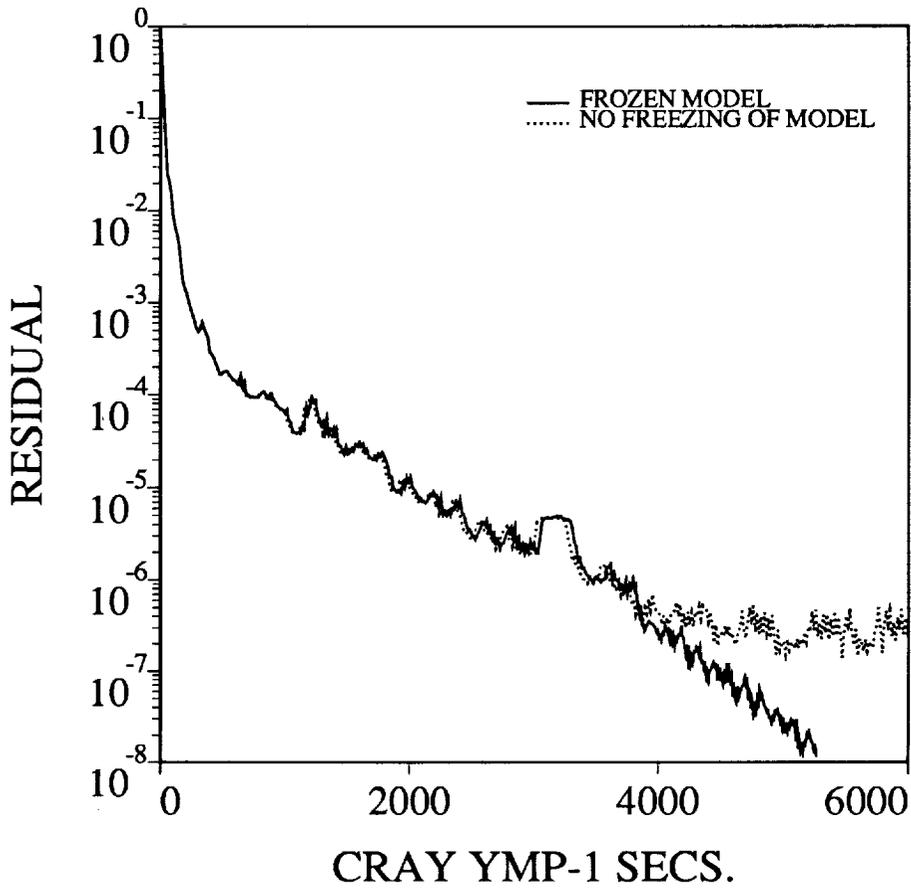


Figure 4c
Convergence History of ILU-GMRES Implicit Scheme and Effect of Freezing
the Algebraic Turbulence Model for Flow Over Four-Element Airfoil
(Mach = 0.1995, Incidence = 16.02 degrees, Reynolds Number = 1.187 million)

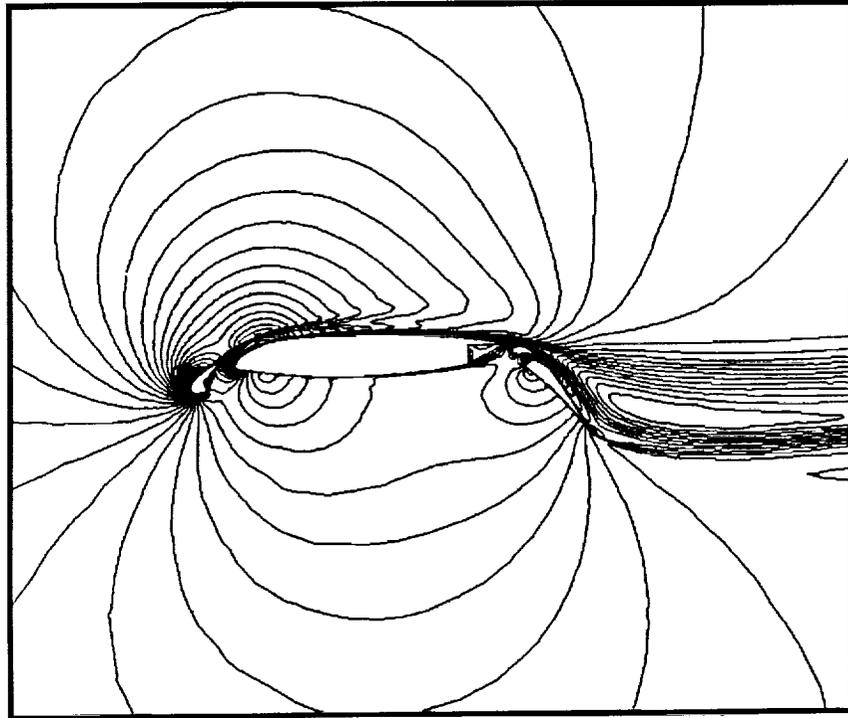


Figure 4d
Computed Mach Contours for Viscous Turbulent Flow Over Four-Element Airfoil
(Mach = 0.1995, Incidence = 16.02 degrees, Reynolds Number = 1.187 million)

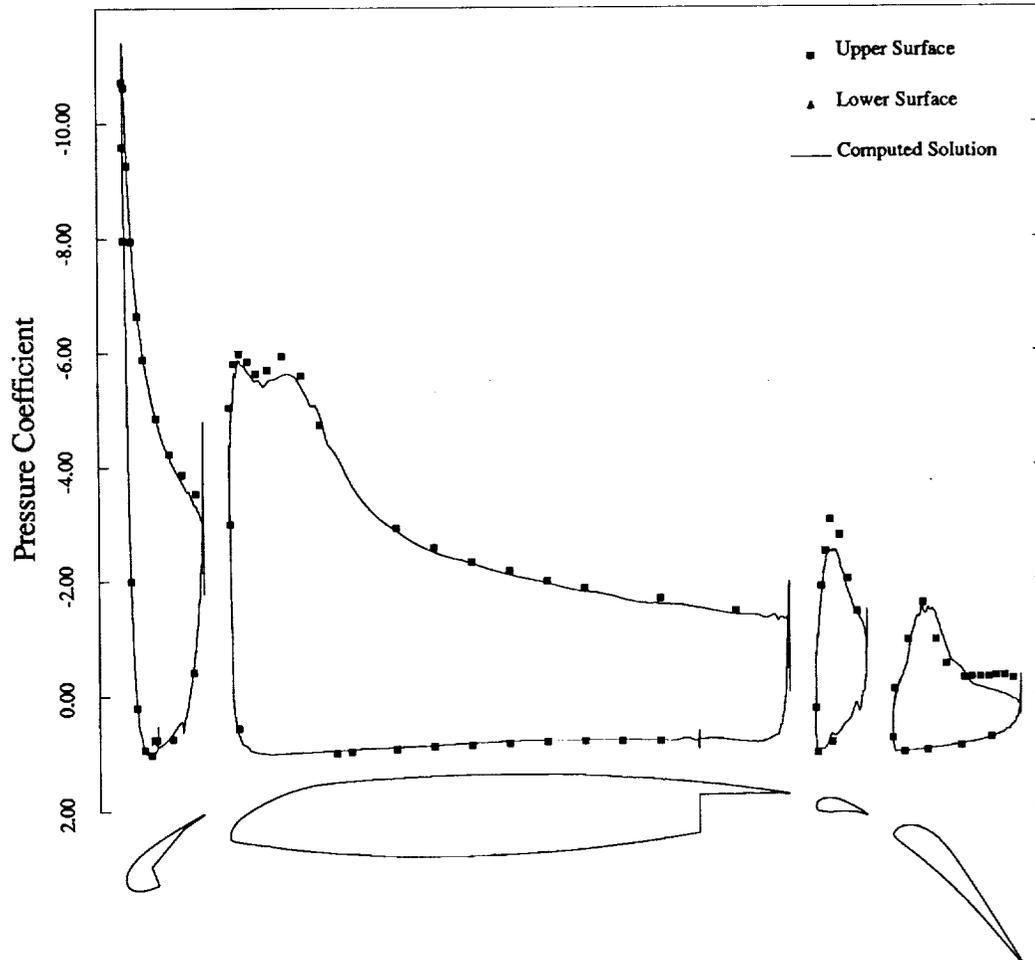


Figure 4e
Computed Surface Pressure Distribution and Comparison with Wind-Tunnel Data
for Viscous Turbulent Flow Over Four-Element Airfoil
(Mach = 0.1995, Incidence = 16.02 degrees, Reynolds Number = 1.187 million)



Report Documentation Page

1. Report No. NASA CR-187564 ICASE Report No. 91-40		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle IMPLICIT SOLVERS FOR UNSTRUCTURED MESHES				5. Report Date May 1991	
				6. Performing Organization Code	
7. Author(s) V. Venkatakrishnan Dimitri J. Mavriplis				8. Performing Organization Report No. 91-40	
				10. Work Unit No. 505-90-52-01	
9. Performing Organization Name and Address Institute for Computer Applications in Science and Engineering Mail Stop 132C, NASA Langley Research Center Hampton, VA 23665-5225				11. Contract or Grant No. NAS1-18605	
				13. Type of Report and Period Covered Contractor Report	
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Langley Research Center Hampton, VA 23665-5225				14. Sponsoring Agency Code	
15. Supplementary Notes Langley Technical Monitor: Michael F. Card Final Report					
16. Abstract <p>We develop and test implicit methods for unstructured mesh computations. The approximate system which arises from the Newton-linearization of the nonlinear evolution operator is solved by using the preconditioned CMRES (Generalized Minimum Residual) technique. We investigate three different preconditioners, namely, the incomplete LU factorization (ILU), block diagonal factorization and the symmetric successive over-relaxation (SSOR). The preconditioners have been optimized to have good vectorization properties. We also study SSOR and ILU themselves as iterative schemes. The various methods are compared over a wide range of problems. Ordering of the unknowns, which affects the convergence of these sparse matrix iterative methods, is also investigated. Results are presented for inviscid and turbulent viscous calculations on single and multi-element airfoil configurations using globally and adaptively generated meshes.</p>					
17. Key Words (Suggested by Author(s)) implicit, unstructured			18. Distribution Statement 02 - Aerodynamics 64 - Numerical Analysis Unclassified - Unlimited		
19. Security Classif. (of this report) Unclassified		20. Security Classif. (of this page) Unclassified		21. No. of pages 23	22. Price A03